Noname manuscript No.
(will be inserted by the editor)

Designing Privacy-aware Internet of Things Applications

Charith Perera · Arosha K. Bandara · Blaine Price · Bashar Nuseibeh

arXiv:1703.03892v1 [cs.SE] 11 Mar 2017

the date of receipt and acceptance should be inserted later

Abstract Internet of Things (IoT) applications typically collect and analyse personal data that can be used to derive sensitive information about individuals. However, thus far, privacy concerns have not been explicitly considered in software engineering processes when designing IoT applications. In this paper, we explore how a Privacy-by-Design (PbD) framework, formulated as a set of guidelines, can help software engineers to design privacy-aware IoT applications. We studied the utility of our proposed PbD framework by studying how software engineers use it to design IoT applications. This user study highlighted the benefits of providing a framework that helps software engineers explicitly consider privacy for IoT applications and also surfaced a number of challenges associated with our approach.

Keywords Internet of Things · Privacy · Software Engineering

1 Introduction

The Internet of Things (IoT) [1] is a network of physical objects or *'things'* that have computing, networking, sensing and actuation capabilities, enabling these objects to collect and exchange data. The design and development process for IoT applications is more complicated than that for desktop, mobile, or web applications. First, IoT applications require both software and hardware (e.g., sensors and actuators) to work together across multiple different type of nodes (e.g., micro-controllers, system-on-chips, mobile phones, miniaturized single-board computers, cloud platforms) with different capabilities under different conditions [2]. Secondly, IoT applications development requires different types of software engineers to work together (e.g., embedded, mobile, web, desktop). This complexity of different software engineering specialists collaborating to combine different types of hardware and software is compounded by the lack of integrated development stacks that support the engineering of end to end IoT applications.

Typically, IoT applications collect and analyse personal data that can be used to derive sensitive information about individuals. However, thus far, privacy concerns have not been explicitly considered in software engineering processes when designing and developing IoT applications, partly due to a lack of Privacy-by-Design (PbD) methods for the IoT. Further, the engineering complexities explained above have forced software engineers to put most of their efforts towards addressing other challenges such as interoperability and modifiability, resulting in privacy concerns being largely overlooked.

We propose to address this issue by providing systematic guidance for software engineers towards building privacy-aware IoT applications.

Our research is motivated by the lack of privacy protection measures taken by IoT application designers and potential privacy violations that arise due to such inefficient designs. In our earlier work [3], we derived privacy guidelines by examining Hoepman's [4] eight design strategies and used them to ***'assess'*** privacy capabilities of IoT applications and platforms. In contrast, our objective in this paper is to explore how a PbD framework (a set of guidelines) can help software engineers to ***'design'*** privacy-aware IoT applications.

The paper is structured as follows: Section 2 discusses common IoT architectures and their characteris-

C Perera · A K. Bandara · B Price · B Nuseibeh
The Open University, Walton Hall,
Milton Keynes, MK7 6AA, UK
Lero, University of Limerick, Ireland
E-mail: firstname.lastname@open.ac.uk

tics. It also briefly introduces the data life cycle phases and their importance when designing privacy into IoT applications. In Section 3, we present our motivation through three different use-cases. We have used these use-cases to evaluate the effectiveness and identify the challenges in designing privacy aware IoT applications. We briefly introduce the PbD framework in Section 4. In Section 5, we explain the research methodology and evaluate the effectiveness the PbD framework. We discuss findings and lessons learned in Section 6. Finally, Section 7 presents the related work and compares our PbD framework with existing approaches followed by the conclusion. In Section 8, we conclude the paper by highlighting future directions.

2 Internet of Things Software Architecture

In this section, we briefly discuss how data flows in a typical IoT application. As illustrated in Figure 1, in IoT applications, data moves from sensing devices to gateway devices to the cloud infrastructure [2]. This is the most common architecture, also called the centralised architecture, used in IoT application development [5]. However, there are other types of architecture such as 1) collaborative, 2) connected intra-net of Things, and 3) distributed IoT [5]. Even for these other types of architectures, if we consider a flow of a single data item, we can observe a data flow analogous to that of the centralised architecture where data moves from edge devices to the cloud through different types of nodes. Therefore, while we use the centralised IoT architecture to explain our PbD approach in this paper, our approach is agnostic the choice of IoT architecture.

The centralised architectures typically consists of three components: 1) IoT devices, 2) Gateway devices, and 3) IoT cloud platforms (Figure 1). Each of these devices have different computational capabilities. They also have different types of access to energy sources from permanent to solar power to battery power. Further, depending on the availability of knowledge, each device may have limitations as to the type of data processing that can be done. A typical IoT application would integrate all these different types of devices with different capabilities. It is important to note that different types of privacy protecting measures can be taken on each of these different nodes based on their characteristics.

We divided the data life cycle into five phases in order to structure our discussion. Data life cycle phases play a major role in applying our PbD framework to design privacy aware applications in a systematic manner. Within each device (also called a node), data moves through five data life cycle phases: Consent and Data

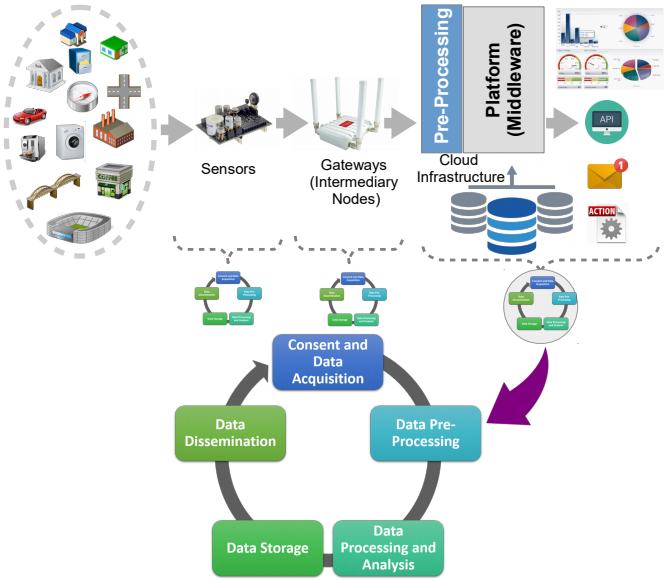

Fig. 1: Typical data flow in Internet of Things Applications

Acquisition [CDA], Data Preprocessing [DPP], Data Processing and Analysis [DPA], Data Storage [DS] and Data Dissemination [DD]. CDA phase comprises routing and data read activities by a given node. DPP describes any type of processing performed on raw data to prepare it for another processing procedure [6]. DPA is, broadly, the collection and manipulation of data items to produce meaningful information [7]. DD is the distribution or transmission of data to an external party.

We assume that all the data life cycle phases are present in all nodes in an IoT application to be utilised by engineers to protect user privacy. However, based on the decisions taken by software engineers, some data life cycle phases in some nodes may not be utilised. For example, a sensor node may utilise the DPP phase to average temperature data. Then, without using either the DPA and DS phases to analyse or store data (due to hardware and energy constraints) the sensor node may push the averaged data to the gateway node in the DD phase.

3 Example IoT Scenarios

In this section, we present three use case scenarios, which we also use to evaluate the PbD framework as described in Section 5. Each scenario is presented from a problem owner's perspective, where each problem could be solved by developing an IoT application. More importantly, it should be noted that none of these scenarios explicitly highlight privacy requirements or challenges. They are primarily focused on explaining func-

tional requirements at a high level. Later in Section 4, we explain how our PbD framework can be used by software engineers to extract additional information, from problem owners, that are crucial to design privacy aware IoT applications.

3.1 Use case 1: Rehabilitation and Recovery

Summary: Robert is a researcher who oversees a number of rehabilitation facilities around the country where patients with physical disabilities are treated and rehabilitated. Robert is interested in collecting and analysing data from sensors worn by patients while they engage in certain activities (e.g., walk using walker, walk using crutches, climbing stairs), in order to guide the patients' recovery processes in a more personalised manner. Robert has an application that is capable of analysing patient data and developing personalised rehabilitation plans. The application monitors the progress and alters the rehabilitation plans accordingly. There is a speciality nurse allocated for each patient in order to monitor the recovery progress and provide necessary advise when required.

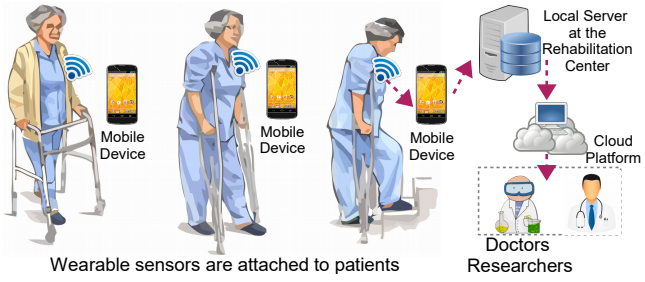

Fig. 2: IoT application to support rehabilitation

3.2 Use case 2: Health and Well-being

Summary: Michael works for the department of public health and well-being. He has been asked to develop a plan to improve the public health in his city by improving the infrastructure that supports exercise and recreational activities (e.g., parks and the paths that supports jogging, cycling, and place for bar exercise, etc.). In order to develop an efficient and effective plan, Michael needs to understand movements of people and several other aspects of their activities. Michael is planning recruit volunteers in order to gather data using sensors. Michael has an application that is capable of analysing different types of data and recommending possible improvements that need to be done. Michael only needs to collect data when the volunteers are within the park premises as illustrated in Figure 3

Fig. 3: City planning towards health and well-being

3.3 Use case 3: Amusement Park and Leisure

Summary: *TrueLeisure* is a company that operates different types of franchised entertainment attractions. Their amusement parks are located in United States, United Kingdom, and Australia. These amusement parks are fully owned and operated by franchisees. However, *TrueLeisure* continuously monitors and assesses the service qualities and several other aspects in each of the amusement parks. Jane is a data analyst overseeing the quality assessment tasks at *TrueLeisure*. She is responsible for continuously monitoring the service quality parameters. Waiting time is one of the key service quality parameters and is a key contributory factor to customer satisfaction. Local quality assessment teams continuously measure the crowd waiting time of each ride and attraction within their own amusement park. All the visitors use *TrueLeisure*'s theme park mobile app to buy tickets for attractions, further information, tour guide, maps, etc. Jane is interested in the big picture, i.e. she would like to measure the overall waiting time for each ride attraction by combining individual waiting times. Jane will report these measurements to *TrueLeisure* management to guide franchisees on future developments of their theme parks efficiently and effectively.

4 Privacy-by-Design Guidelines

In each of the example scenarios above, the software engineer would need to perform further analysis to extract explicit privacy requirements that could support the design of privacy enhancing features into the IoT applications that would be developed to delivery the required

Table 1: Privacy-by-Design Framework

Guideline	DA	DPP	DPA	DS	DD	Capability	Minimise	Hide	Separate	Aggregate	Inform	Control	Enforce	Demonstrate	Privacy Risks
1-Minimise data acquisition	✓	✓					✓	✓							⊗ ⊖
2-Minimise number of data sources	✓						✓								⊗ ⊖
3-Minimise raw data intake	✓	✓					✓			✓					⊗ ⊖
4-Minimise knowledge discovery			✓				✓								⊗
5-Minimise data storage				✓			✓								⊗ ⊖
6-Minimise data retention period				✓			✓	✓							⊗ ⊖
7-Hidden data routing	✓				✓			✓							⊖
8-Data anonymisation	✓	✓	✓		✓			✓							⊗ ⊖
9-Encrypted data communication	✓				✓			✓							⊖
10-Encrypted data processing		✓	✓					✓							⊖
11-Encrypted data storage								✓							⊖
12-Reduce data granularity	✓	✓	✓		✓										⊗
13-Query answering					✓		✓	✓							⊗
14-Repeated query blocking					✓		✓	✓							⊗
15-Distributed data processing			✓						✓						⊗ ⊖
16-Distributed data storage				✓					✓						⊗ ⊖
17-Knowledge discovery based aggregation	✓	✓	✓	✓	✓					✓					⊗
18-Geography based aggregation	✓	✓	✓	✓	✓					✓					⊗
19-Chain aggregation	✓	✓	✓	✓	✓					✓					⊗
20-Time-Period based aggregation	✓	✓	✓	✓	✓					✓					⊗
21-Category based aggregation	✓	✓	✓	✓	✓					✓					⊗
22-Information Disclosure	✓	✓	✓	✓	✓						✓			✓	⊗
23-Control	✓	✓	✓	✓	✓							✓	✓		⊗
24-Logging	✓	✓	✓	✓	✓									✓	⊗ ⊖
25-Auditing														✓	
26-Open Source														✓	
27-Data Flow Diagrams														✓	
28-Certification														✓	
29-Standardisation														✓	
30-Compliance														✓	⊗ ⊖

Risk Types: Secondary Usage (⊗), Unauthorised Access (⊖)

functionality. In this section we provide an overview of our PbD framework and explain how it could be used to design privacy into IoT applications.

4.1 Overview of the Framework

In our earlier work [3], our literature search led to determine that Hoepman's [4] is the most appropriate starting point for developing a more detailed PbD framework for IoT. Detailed explanation on each of the guidelines and reasoning behind the extraction of each guideline are presented in [3]. However, these guidelines are not fool-proof recommendations that can be used naively. Each IoT application is different in term of their objectives, implementations, execution, etc.

We developed these guidelines to act as a framework to support software engineers, so they can adopt our guidelines into their IoT applications in a customised manner. For example, certain applications will require aggregation of data from different sources to discover a certain new knowledge (i.e. new pieces of information). We do not discourage such approaches as long as data is acquired through proper consent acquisition processes. However, IoT applications, at all times, should take all possible measures to achieve their goals with a minimum amount of data. This means that out of eight pri-

vacy design strategies proposed by Hoepman [4], minimisation is the most important strategy.

The relationship between Hoepman's [4] design strategies and our guidelines are presented in Table 1. At a high-level, we have identified two major privacy risks, namely, secondary usage ($\otimes$) and unauthorized access ($\ominus$) that would arise as consequences of not following the guidelines. Secondary usage refers to the use collected data for purposes that were not initially consented to by the data owners [8], which can lead to privacy violations. Unauthorised access is when someone gains access to data without proper authorisation during any phase of the data life cycle. We will use the above symbols to denote which threat is relevant to each guideline. In Table 1, privacy guidelines are colour coded based on the primary privacy design strategy that they are belong to. However, it is important to note that some guidelines may belong to multiple design strategies. For example, **(Guidelines 6)** *minimise data retention period* can primarily be identified as a minimise strategy, but it can also be classified as a hide strategy as it reduces the period for which data is visible.

4.2 Use of Privacy-by-Design Framework

One of the primary objectives of the proposed PbD framework is to help software engineers to ask the right questions regarding privacy protection when designing IoT applications and their architectures. These guidelines provide them with a framework to start thinking about privacy and direct them to incorporate privacy features into IoT application designs. A piece of software is designed to solve a problem. Sometimes, a problem may be identified by a person who is affected by the problem (e.g., Robert, Michael or Jane in our motivating scenarios). At other times, a third party company may identify a generic problem that affects many other people (e.g., Enterprise resource planning solutions). This type of software engineering is common in the IoT domain as well. Some IoT solutions are generic middleware platforms that can be used to build end to end applications. Others are complete IoT applications that aim to solve a specific problem [2, 9].

However, problem owners mainly focus on the requirements that would help to solve their problem [10], ignoring privacy considerations. Therefore, privacy requirements are largely overlooked when designing software architectures for IoT applications. The PbD framework allows both problem owners and software engineers to sit together and discuss the problem and incorporate privacy protecting measures into IoT application designs.

In section 3, we presented three use case scenarios. For each scenario, we have a problem owner's expectation and a brief set of requirements. There are no explicit reference to privacy protecting measures. We assume, additional information can only be gathered through questioning the problem owners and domain experts. In the user study, we simulated such discussions between the problem owners (i.e, represented by ourselves, the researchers) and the software engineers (i.e., represented by the study participants). Our hypothesis was that the PbD framework helps software engineers to ask questions from both problem owners and domain experts in order to extract detailed requirements that could be used to design privacy into IoT applications.

Let us revisit the scenario presented in section 3.1 in order to demonstrate how to use the PbD framework to extract privacy requirements towards designing a privacy-aware IoT application.

Guideline 1 leads software engineers to ask the question: what type of data is required to achieve the Robert's objective? In our scenario the problem owner responds as follows:

Robert collects data using wearable sensor kits. The collected data types are pulse, oxygen in blood (SPO2), airflow (breathing), body temperature, electrocardiogram (ECG), glucometer, galvanic skin response (GSR-sweating), blood pressure (sphygmomanometer), patient activity (accelerometer) and muscle / eletromyography sensor (EMG). Accelerometer are used to derive patient activity. In addition to the sensor data, weather information such as temperature, humidity are also important for the Robert's research. Patients' mobile phones GPS sensors and weather APIs are used to collect such information. The data collection sampling rate is expected to be 30 seconds. Data is only required to be collected when patients are performing either one of the monitored activities (i.e. walking with walker or crutches, or climbing stairs).

Based on this information the software engineer can decide not acquire any other types of data and also design appropriate sampling rate controls into the application. This will have the effect of minimising data acquisition and reducing the risk of both secondary usage and unauthorised access to private data.

In a similar fashion, guidelines 3, 5, 20 and 21 would lead a software engineer to ask questions such as: what type of data is required in raw format and what type of information can be aggregated in order to reduce privacy risks?. As a result, the following information may be gathered.

Robert requires oxygen in blood (SPO2), airflow (breathing), body temperature data types in raw format and need to be accurate. The data collection sampling rate

is expected to be five seconds. In contrast, other data items can be aggregated into averaged values (e.g., aggregated over two minutes).

Similar guidelines based questioning can be used to extract privacy requirements which the software engineers can use to systematically design privacy aware IoT applications. Due to space limitations, we don't detail all the questions that could be asked in relation to the scenario. Instead, below we provide the information that could be acquired using our PbD approach by annotating a detailed description of the scenario with references to the relevant PbD guidelines at the end of each statement.

*The sensor kit is expected to push data to the patient's mobile phone using Bluetooth. The mobile phone pushes data to the rehabilitation centre's local server using Wi-Fi. The local server pushes data to the cloud IoT platform. Patients come to the rehabilitation centre 3 days a week in order to perform the tasks assigned to them. Another 3 days they perform the task at their homes. The smart phone is expected to push data to the local server at the end of each day (**Guideline 6**). However, if the patients perform their tasks at home, data need to be kept stored on the mobile until the next time they visit the rehabilitation centre (**Guideline 6**). The speciality nurses monitor the progress and advice the patients on weekly basis. The speciality nurses' responsibility is to make sure that the patient are performing the tasks as assigned by the recommendation system and assists patients if they have any difficulties in following the assigned tasks and schedules. Robert is required to analyse data every six months in order to understand the how to improve the rehabilitation processes in a personalized manner (**Guideline 6**). For long term data analysis purposes, Robert's application stores data after averaging over five minute (**Guideline 6**).*

*Robert's application requires average over five minute when patients are performing their tasks (**Guideline 20**). Patient data can be anonymized (**Guideline 8**). Data storages in moth mobile device, local server and Robert's cloud server should store data in encrypted form (**Guideline 11**). End-to-end encryption can be used to secure the data communication (**Guideline 9**). Robert does not require exact location of the patient where they may have performed the activities. The requirement is to acquire the weather parameters such as temperature, humidity, etc. Therefore, location data can be abstracted without affecting the accuracy of the data (**Guideline 12**). In this IoT application, data processing and storing happens in three different locations (nodes), namely, 1) patient phone, 2) local server, and 3) Robert's cloud server (**Guideline 15 and 16**).*

The above example illustrates how our PbD guidelines could be used to extract additional information regarding a use case which enables software engineers to design appropriate privacy enhancing features into their IoT applications. In order to evaluate the effectiveness of our approach, we developed similar detailed requirement descriptions for each of the use case scenarios, which we have omitted here due to space limitations. It is important to note that not all privacy guidelines are relevant to all IoT applications. In Table 2, we summarise which privacy guidelines are relevant to each scenario.

Table 2: Relevant Privacy Requirements for Each Scenario

Guideline (↓) Use Case Number (→)	1	2	3
1-Minimize data acquisition	✓	✓	✓
2-Minimize number of data sources	–	✓	–
3-Minimize raw data intake	✓	✓	✓
5-Minimize data storage	✓	✓	✓
6-Minimize data retention period	✓	✓	✓
7-Hidden data routing	✓	✓	✓
8-Data anonymization	✓	✓	✓
9-Encrypted data communication	✓	✓	✓
11-Encrypted data storage	✓	✓	✓
12-Reduce data granularity	✓	✓	✓
15-Distributed data processing	✓	✓	✓
16-Distributed data storage	✓	✓	✓
18-Geography based aggregation	–	–	✓
20-Time-Period based aggregation	✓	✓	✓
21-Category based aggregation	✓	✓	✓
	13	14	14

5 Evaluation

This section explains how we conducted the evaluations and justifies our research methodology. First, we will explain the aims of this study. Next, we explain and justify the participant recruitment strategy and sample size before describing the procedures followed at each step of the study.

5.1 Purpose

The purpose of this work is to explore how our PbD framework can help software engineers to design privacy-aware IoT applications. Through user studies, using quantitative and qualitative data analysis, we aimed to answer following three questions that explore the effectiveness of the proposed PbD framework. We discuss these questions later in this section.

– Can the proposed PbD framework guide novice software engineers to design IoT applications that are more privacy-aware than they would do otherwise?

– Can the proposed PbD framework guide expert software engineers to design IoT applications that are more privacy-aware than they would do otherwise?
– Out of novice and expert software engineers, who would benefit most from the proposed PbD framework? or in other words, does the software engineering expertise matter when it comes to incorporating privacy protection features into IoT application designs?

In the first two questions above, we consider the design of a IoT application to be more privacy-aware if it considers a greater number of privacy concerns to incorporate appropriate privacy protecting features. We measure this in terms of the number of privacy guidelines considered by the study participants when designing the example IoT applications.

5.2 Recruitment and Remuneration

In total, we recruited 10 participants for the study of which five were novice software engineers and five were expert software engineers. A participant was classified as a novice if they had less than three years of experience (full time) in a software engineering role (design or development). Participants with more than three years of experience (design or development), were considered to be experts. We adopted an opportunistic sampling technique and recruitment was from the staff and student populations at The Open University and the University of Surrey. No criteria other than software engineering experience was considered when recruiting. We collected demographic information such as age, highest education qualification, and the number of years in a software engineering role. Each participant was compensated with shopping vouchers valued at GBP 20. There were no failure criteria as long as the participant attend the data collection session of the study. The study design was reviewed and approved by the Human Research Ethics Committee of The Open University.

5.3 Procedure

All the data collection sessions were carried out as 1-to-1 lab-based observational studies [11]. The principal investigator (PI) acted as the facilitator as well as the observer during each of the sessions. The duration for each session was 1.5 hours. At the beginning of the each session, participants were given the consent from to sign off and brief demographic information were collected. We audio recorded all the discussions between the participants and the PI for qualitative analysis purposes. Next, participants were given a instruction sheet, as shown in Figure 4, that comprised a set of example notations that could be used to illustrate the design of the IoT applications. Participants were reassured that adherence to the notation was not essential.

We divided the rest of the study into three rounds, which we call **Round 1 (NoPrivacy):** *IoT application design without any guidance to consider privacy or reference to the PbD guidelines*, **Round 2 (WithPrivacy):** *IoT application design without privacy guidelines but guidance to consider privacy*, **Round 3 (WithPbDGuidelines):** *IoT applications design with privacy guidelines*. However, this segmentation was only used to structure the discussions and observations and none of them were formally acknowledged or identified during the interviews.

Round 1: It is important to note that we only informed the participants that this is an IoT application design study, without making any reference to privacy. By doing so, we expected them to be unbiased and follow the steps they would have naturally followed in designing an IoT application. We gave them separate A4 sheets to draw their IoT application designs with respect to each use case. They were briefed about the notations they could use, but we did not restrict them to any particular notation as long as their designs are understandable and clearly annotated.

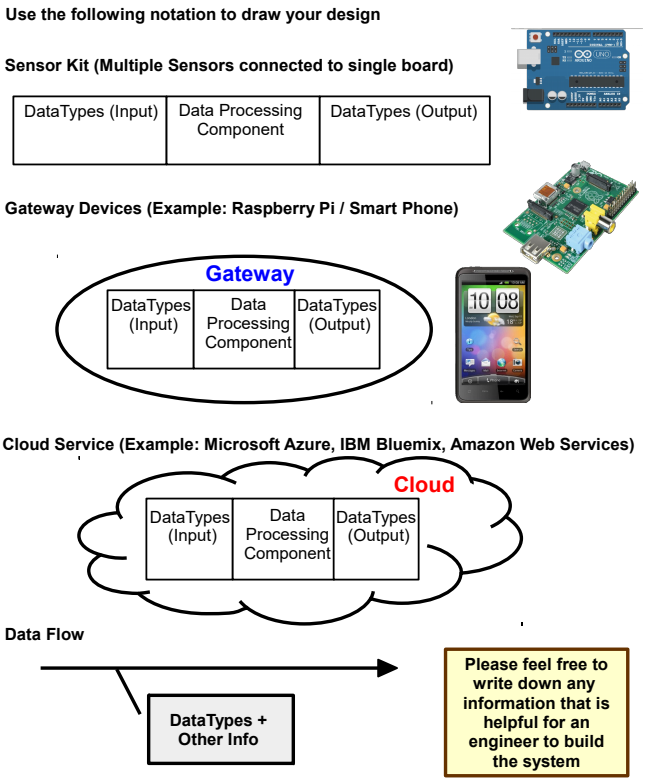
Use the following notation to draw your design
Sensor Kit (Multiple Sensors connected to single board)
DataTypes (Input)
Data Processing Component
DataTypes (Output)
Gateway Devices (Example: Raspberry Pi / Smart Phone)
Gateway
DataTypes (Input)
Data Processing Component
DataTypes (Output)
Cloud Service (Example: Microsoft Azure, IBM Bluemix, Amazon Web Services)
Cloud
DataTypes (Input)
Data Processing Component
DataTypes (Output)
Data Flow
DataTypes + Other Info
Please feel free to write down any information that is helpful for an engineer to build the system

Fig. 4: Notations to be used in IoT application Design

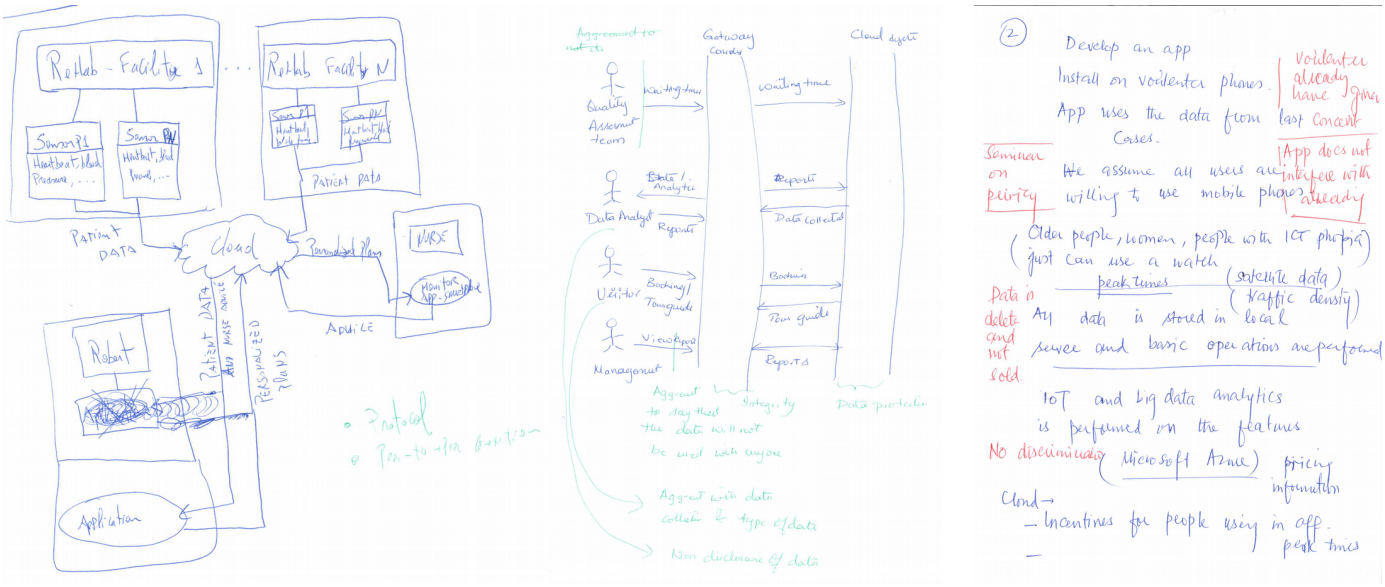

Fig. 5: Sample IoT application designs that illustrate a variety of approaches used by participants to express their high-level designs. In addition to block diagram notations based on the examples we provided, participants used sequence diagrams, pictorial diagrams and detailed text descriptions as illustrated above.

Next, the participants and they were asked to design IoT applications to satisfy the requirements of each the scenarios presented in Section 3. Initially the participants worked from the summary descriptions provided in this paper but the PI was prepared to provide more detailed information, similar to that presented in Section 4.2 if the participant explicitly asked any related questions. We designed the study to simulate a conversation between a software engineer and a problem owner where the engineer is trying to elaborate the requirements and design the architecture of the IoT application.

We encouraged participants to ask as many question as possible about the case studies and application requirements. This means that participants could have asked any question regarding privacy requirements if they wanted to. Some of the commonly asked questions are discussed later in this paper. We gave them 50 minutes to complete the IoT application designs for the three use cases provided. However, the time limit was given as a guidance and we did not enforce it. The actual time of each study was varied based the the amount taken by the participants on each phase. So the actual total time varied between 1 hour and 15 minutes to 2 hours. We always allowed each participant to naturally progress through their designs without rushing them through each phase. After the designs were completed, we asked the participants to explain their designs and briefly justify their design decisions.

Round 2: Next, we gave participants a ten minute introduction on privacy. In order to achieve consistency, accuracy, and a well recognised description of privacy and related challenges, we selected two videos[1] [2] from You Tube produced and published by *Privacy International* (www.privacyinternational.org). The objective of showing these videos to each participant was to provoke them to think about privacy and help them to recall their past experiences and knowledge of dealing with privacy issues. This was intended to help them with the next task. It is important to note that we did not provide any additional material on privacy at this stage.

Next, we asked the participants to refine their previous IoT application designs further to protect user privacy. Similar to the previous round, questions were welcomed. We gave the participants 20 minutes to refine the IoT application designs for the three use cases provided. For Round 2, they wrote in a different colour to round 1, which enabled us to distinguish the design activities from each round clearly. After the revisions were made, we asked the participants to explain their revised designs and how they improved the privacy protection.

Round 3: Finally, we gave participants an introduction on the PbD guidelines and how to use them. We asked the participants to refine their previous IoT application designs to protects user privacy. Similar to the

[1] What Is Privacy? (youtube.com/watch?v=zsboDBMq6vo)
[2] Big Data (youtube.com/watch?v=HOoKhnvoYkU)

previous round, questions were welcomed. We gave the participants 20 minutes to enhance the privacy features of their IoT application designs for the three use cases provided. After the revisions were made, we asked the participants to explain their revised designs and how they improve the privacy protection. Once completed, we collected the IoT application designs produced by the participant.

Some sample application designs produced by participants are presented in Figure 5.

6 Findings, Discussion and Lessons Learned

In this section, first, we analyse and discuss the results quantitatively. Our aim is to address the three questions presented earlier in Section 5 with the help of data collected through the users studies. Later, we discuss the results qualitatively in order to understand software engineers' approach towards designing privacy-aware IoT applications.

6.1 Exploring Effectiveness (Quantitative Analysis)

As shown in Table 2, we expected each participant to identify a maximum of 41 privacy protecting measures (Use-case 1: 12 measures, Use-case 2: 14 measures, Use-case 3: 14 measures). The participants may identify these privacy measures either using their experience, common sense, or using the PbD guidelines. In total, we collected 410 data points (41 measures x 10 participants). We have presented a snapshot view of the data gathered using two heat-maps in Figure 6 where the results for novice and expert software engineers are presented separately.

The heat-maps clearly show that both novice and expert software engineers were able to identify a greater number of privacy protecting measures by using the PbD guidelines than they would do otherwise. In Figure 7, we illustrate how the mean of the *'number of privacy measures'* identified, by both novice and experts software engineers, changes at different privacy knowledge levels. The average number of privacy measures identified, in Round 1, by novices is 0.2 and experts is 2.2. Similarly, the average number of privacy measures identified, in Round 2, by novices is 6.6 and experts is 6.8. Further, the average number of privacy measures identified, in Round 3, by novices is 32.6 and experts is 30.4.

Next, we ran statistical tests (i.e., ANOVA[3]) and found out that there is a significant difference between

[3] statistics.laerd.com/statistical-guides/one-way-anova-statistical-guide.php

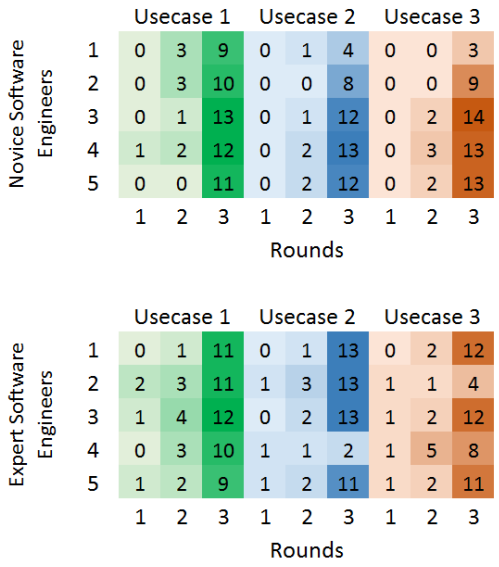

Fig. 6: The three use-cases are marked using three separate colours. The x-axis denotes how much privacy protecting measures have been identified in each round (higher the number of privacy requirements identified, darker the shading is). The y-axis denotes the participant ID.

the number of privacy measures identified with and without the PbD guidelines (within=PrivacyKnowledge (ANOVA p = 2.099781e-09; p < 0.05)). Further, our results show that the expertise of the software engineers (novice vs. expert) has no significant effect on the ability of identifying privacy protecting measures (between=Expertise (ANOVA p = 6.897806e-01; p < 0.05)) .

Figure 8 illustrates which privacy guidelines have been identified in each round by the participants. It

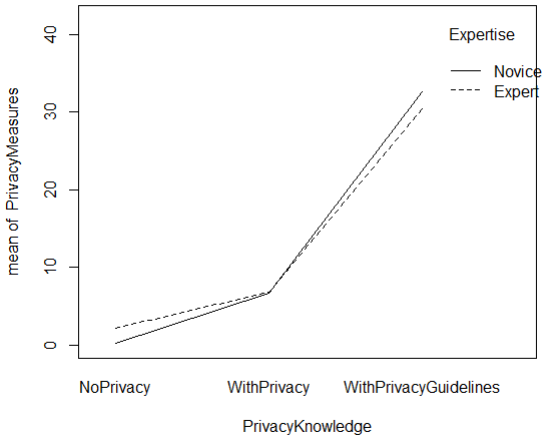

Fig. 7: Number of privacy measures identified in each round

is also important to note that PbD guideline 2 and 18 were only relevant in one of the use case scenarios (out of three) which justifies its unusually low identification rate. Comparatively, more participants have identified PbD guideline 3 (Minimise raw data intake) and 20 (Time period based aggregation) in Round 1. However, our discussions revealed that participants integrated these features into their designs to meet functional requirements of the scenarios rather than due to a consideration of privacy. In Round 2, after we explicitly asked them to improve the privacy awareness of their IoT application designs, participants primarily identified guidelines 8 (data anonymisation), 9 (encrypted data communication), and 11 (encrypted data storage). In Round 3, there was no noticeable different in the guidelines identified by the participants.

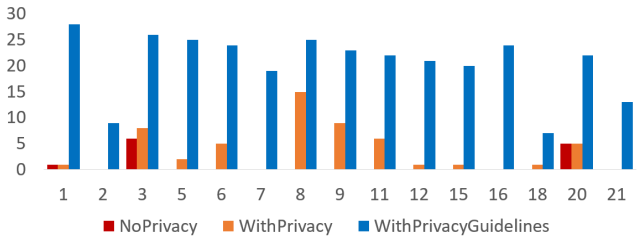

Fig. 8: Privacy guidelines identified in each round: the x-axis denotes privacy guidelines by number and each colour represents the three rounds. The y-axis denotes the frequency which participants identified a given privacy guideline.

In total, we expected participants to identify a maximum of 410 privacy preserving measures that they could take in order to improve the privacy awareness of the three given IoT application scenarios. They identified 308 privacy preserving measures with the help of the PbD guidelines. Therefore, the success rate is 75.12%. As shown in Figure 7, this results is significantly better than *'without PbD guidelines'*. Based on our discussions with the participants, we identified two main reasons why sometimes they failed to apply a given guideline into their application designs: 1) unique IoT application designs eliminates the necessity of applying certain privacy preserving measures and 2) the lack of time. It is important to note that the PbD guidelines can only be applied to protect user privacy in certain application design contexts. Some participants designed their IoT applications in such away that certain PbD guidelines have no role to play. We discuss one such example in the next section.

6.2 Exploring Effectiveness (Qualitative Analysis)

Our study showed that software engineers do not consider privacy as a first class citizen in their IoT application designs. This justifies our decision to develop a PbD framework to guide the thought process of software engineers. During our user studies, participants candidly expressed their wish to collect as much data as possible (e.g., Participants E2 said *"As a developer, we like all of the data"*). We believe that this mindset of collecting as much data as possible needs to be changed towards a *Privacy Mindset* where only the most essential data items are gathered and processed. We explained the privacy risks of gathering non-essential data in Section 4.

We also observed post hoc rationalisation from most of the participants. After we showed the PbD guidelines, most of the participants felt the responsibility of addressing privacy issues in their IoT application designs. Most of them not only followed the guidelines and successfully improved their designs, but also claimed that they thought about certain privacy considerations before we showed them the guidelines, even though their designs did not show any evidence of this. This behaviour suggests that software engineers are well aware of the importance of privacy issues, though they do not make any effort to address them until an external impetus (e.g., we as researchers in this case or it could be pressure from laws and regulations, or explicit client demands) that explicitly encourage them to do so. When we explicitly encouraged them to address privacy issues, most of the participants felt the need of defending themselves and claim that they thought about privacy before. This post hoc rationalisation behaviour justifies the importance of developing a *Privacy Mindset* among software engineers.

An important aspect of IoT application design is the choice of the right sensors and techniques to collect data. We realised that these choices also have a direct impact on the privacy. In relation to Scenario 2 (section 3.2), one of our participants (E4) used stationary sensors that do not capture any personally identifiable information to collect the necessary data (e.g., pressure sensors deployed in the ground, motion sensors, infrared sensors, and so on). Sensor technologies have their own strengths and weakness. Similarly, privacy risks also varies depending on the technology used. However, the decision on which technology to use is based on the exact application, cost associated, and how much privacy risks willing to take. For example, deploying pressure sensors on different paths of a given park would eliminate the necessity of hiring volunteers with wearable sensor kits and associated privacy risks. However,

deploying such sensing technology in real world could be much more challenging, in terms of cost, time, and effort, than distributing number of sensor kits among volunteers. On the other hand, stationary sensors would eliminate the hassle of recruiting volunteers, managing them, and their sensor kits. The lesson is that privacy risks can also be reduced by selecting certain types of sensing technologies given that they are feasible to use in a given IoT application.

It is also important to note that three participants identified authentication as a measure of protecting user privacy. However, in our PbD framework we considered authentication as to be a security measure rather than a privacy protection measure. Further, three participants highlighted the importance of acquiring consent from data owners before collecting data. They also pointed out the importance of giving control to the data owners so they can decide on which data to share. Both consent acquisition (information disclosure - guidelines 22) and control (guidelines 23) appeach in our PbD framework even though we did not use them in this evaluation.

6.3 Limitations

Although all the participants were able to understand our proposed guidelines, it was apparent that familiarity is key to applying them in a given IoT application design in a short period of time. For our study, we printed the PbD guidelines on plain A4 sheets as a list. However, the experience of our study participants highlighted that this type of printed list is difficult to follow and can be more time consuming to use. We believe that approaches such as *Privacy Ideation Cards* [12] and KnowCards[4] would be more effective by allowing users to quickly familiarise themselves with the guidelines. In particular, using a colour coded, iconographic approach to represent the guidelines could improve the users recall them and thus lead to faster application of guidelines with less frustration.

An additional limitation of this work is that we did not consider the adaptive nature of privacy. While some decisions about implementing privacy preserving measures can be taken at design-time, IoT applications are by nature unpredictable. As a result, the ability to adapt is an important feature in IoT applications. Ideally, IoT applications should be able to compose built-in privacy preserving techniques into a run-time configuration, that maximises the privacy protection level while maintaining the overall utility of the application.

[4] know-cards.myshopify.com

7 Related Work

There are a number of existing frameworks that have been proposed to help elicit privacy requirements and to design privacy capabilities in systems. The original PbD is a framework proposed by Ann Cavoukian [13], the former Information and Privacy Commissioner of Ontario, Canada. This framework identifies seven foundation principles by which privacy sensitive application should be developed. These are: (1) proactive not reactive; preventative not remedial, (2) privacy as the default setting, (3) privacy embedded into design, (4) full functionality positive-sum, not zero-sum, (5) end-to-end security-full life-cycle protection, (6) visibility and transparency- keep it open, and (7) respect for user privacy, keep it user-centric. The ISO 29100 Privacy framework [14] has proposed eleven design principles, namely, (1) consent and choice, (2) purpose legitimacy and specification, (3) collection limitation, (4) data minimisation, (5) use, retention and disclosure limitation, (6) accuracy and quality, (7) openness, transparency and notice, (8) individual participation and access, (9) accountability, (10) information security, and (11) privacy compliance. In similar direction, Luger et al. [12] aims to understand how to make emerging European data protection regulations more accessible to general public using a series of privacy ideation cards. They have extracted 40 design principles by examining EU General Data Protection Regulation 2012 Com Final 11 [15]. These high level principles are proposed for computer systems in general but not prescriptive enough to be adopted by IT professionals for designing and developing IoT applications.

Building on the ideas of engineering privacy by architecture vs. privacy-by-policy presented by Spiekerman and Cranor [16], Hoepman [4] proposes an approach that identifies eight specific privacy design strategies: minimise, hide, separate, aggregate, inform, control, enforce, and demonstrate. This is in contrast to other approaches that we considered. In a similar vein, LINDDUN [17] is a privacy threat analysis framework that uses data flow diagrams (DFD) to identify privacy threats. LINDDUN focuses on eliminating set of pre-identified privacy threats using a systematic review of data flow diagrams. It consists of six specific methodological steps: (1) define the DFD, (2) map privacy threats to DFD elements, (3) identify threat scenarios, (4) prioritize threats, (5) elicit mitigation strategies, and (6) select corresponding privacy enhancing technologies. However, both LINDDUN and Hoepman's framework are not aimed at the IoT domain. Further, they not prescriptive enough in guiding software engineers.

In contrast, the STRIDE [18] framework was developed to help software engineers consider security threats, is an example framework that has been successfully used to build secure software systems by industry. It suggests six different threat categories: (1) spoofing of user identity, (2) tampering, (3) repudiation, (4) information disclosure (privacy breach or data leak), (5) denial of service, and (6) elevation of privilege. However, its focus is mostly on security than privacy concerns.

8 Conclusions and Future Work

In this paper, we explored how a PbD framework, formulated as a set of guidelines, can help software engineers to design privacy-aware IoT applications. We evaluated the effectiveness of the proposed PbD framework through a use cases based observational study where the participants were asked to design IoT applications to satisfy three given use cases. According to our findings, the proposed PbD guidelines framework has significantly improved the privacy awareness of the IoT applications designed by both novice and expert users. Further, results show that software engineering expertise does not matter significantly when it comes to incorporating privacy protection features into IoT application designs.

In the future, we will conduct research to develop a set of privacy tactics and patterns that are less abstract than guidelines. Such tactics and patterns will help software engineers to tackle specific privacy design challenges in IoT domain. At the moment, privacy guidelines are presented to the software engineers in plain text organised into a list. Though it is usable, in the future, we will explore how we can use human computer interaction (HCI) techniques to make these PbD guidelines more user friendly and accessible to the software engineers. HCI techniques will help software engineers to efficiently and effectively browse and find relevant privacy guidelines, patterns and tactics in a given IoT application design context.

Acknowledgement

We acknowledge the financial support of European Research Council Advanced Grant 291652 (ASAP).

References

1. C. Perera, A. Zaslavsky, P. Christen, D. Georgakopoulos, Communications Surveys Tutorials, IEEE **16**(1), 414 (2014)
2. C. Perera, C.H. Liu, S. Jayawardena, IEEE Transactions on Emerging Topics in Computing **3**(4), 585 (2015)
3. C. Perera, C. Mccormick, A. Bandara, B.A. Price, B. Nuseibeh, in *The 6th International Conference on the Internet of Things* (2016), pp. 83–92
4. J.H. Hoepman, in *ICT Systems Security and Privacy Protection*, *IFIP Advances in Information and Communication Technology*, vol. 428, ed. by N. Cuppens-Boulahia, F. Cuppens, S. Jajodia, A. Abou El Kalam, T. Sans (Springer Berlin Heidelberg, 2014), pp. 446–459
5. R. Roman, J. Zhou, J. Lopez, Computer Networks **57**(10), 2266 (2013)
6. D. Pyle, *Data preparation for data mining* (Morgan Kaufmann Publishers, San Francisco, Calif, 1999)
7. C.S. French, *Data Processing and Information Technology* (Cengage Learning Business Press, 1996)
8. W. Lowrance, Journal of Health Services Research & Policy **8**(suppl 1), 2 (2003)
9. C. Perera, C.H. Liu, S. Jayawardena, IEEE Access **2**, 1660 (2014). DOI 10.1109/ACCESS.2015.2389854
10. L. Bass, P. Clements, R. Kazman, *Software Architecture in Practice*, 3rd edn. (Addison-Wesley Professional, Upper Saddle River, NJ, 2012)
11. H. Sharp, Y. Rogers, J. Preece, Book **11**, 602 (2011). DOI 10.1162/leon.2005.38.5.401
12. E. Luger, L. Urquhart, T. Rodden, M. Golembewski, Proceedings of the ACM CHI'15 Conference on Human Factors in Computing Systems **1**, 457 (2015)
13. A. Cavoukian, in *32nd International Conference of Data Protection and Privacy Commissioners* (2010)
14. ISO/IEC 29100, Information technology Security techniques Privacy framework. Tech. rep. (2011)
15. European Commission, Proposal for a Regulation of the European Parliament And Of The Council on the protection of individuals with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation). Tech. rep. (2012)
16. S. Spiekermann, L. Cranor, IEEE Transactions on Software Engineering **35**(1), 67 (2009)
17. M. Deng, K. Wuyts, R. Scandariato, B. Preneel, W. Joosen, Requirements Engineering **16**(1), 3 (2011)
18. M. Howard, S. Lipner, *The security development lifecycle: SDL, a process for developing demonstrably more secure software* (Microsoft Press, 2006)